Drive technology overview

technology brief, 3rd edition

Abstract 2

Introduction 2

Categories of server disk drives 2

Characteristics of disk drives 3
- Small form factor and large form factor disk drives 3
- Disk drive capacity 4
- Disk drive performance 4
- Reliability 5

Drive qualification process 6

Interconnect technology 7
- Serial Attached SCSI (SAS) 7
- SAS-2 8
- Second generation Serial ATA (SATA) 8
- Interconnect bandwidths and drive throughput 8

Improved performance and reliability with RAID 9

Advanced controllers 9

Solid state drives for servers 9
- HP solid state drive reliability requirements 10
  - NAND over provisioning and wear-levelling 10
  - Power loss protection 10
- Value, mainstream and performance SSDs 11
- SSD performance 12

Conclusion 12

For more information 13

Call to action 13

Downloaded from www.Manualslib.com manuals search engine

Abstract

Disk drive capacity has increased at rates similar to those of microprocessor performance. New and improved interconnect technologies allow the rapid transfer of large amounts of data to and from the disk. New and more cost-effective applications are made possible by increased storage capacity and reduced cost per bit. Through innovation and engineering expertise, HP develops industry-leading disk technologies that optimize overall system capacity, performance, reliability, and value.

This technology brief reviews the classes of disk drives, the key factors determining capacity, performance and reliability in single drives, the options available to connect the drives to the system, and the use of multiple drives to further increase performance and reliability.

Introduction

Disk drives provide the primary mechanism for storing and retrieving permanent, or non-volatile, data in almost all servers, desktop computers, and notebook computers. Disk drives are also increasingly common in portable electronic devices such as music players and automobile navigation systems.

The key performance differences between main memory (semiconductor RAM) and primary storage (typically, magnetic disk drives) are speed of access and capacity. Accessing a disk drive is typically approximately 100,000 times slower than accessing main memory. Primary storage is typically at least 100 times larger than main memory. Most engineering in disk drive and interconnect technologies is driven by a desire to simultaneously reduce the difference in access speed and increase disk drive capacity and reliability. Innovative strategies in disk and disk controller design continue to deliver dramatic increases in disk capacity, performance, and availability.

Flash memory technology, which has previously been used as a low-performance, lower capacity storage medium in consumer devices, is being adapted as primary storage in computers. This technology has the promise of delivering enterprise class storage with lower latencies and significantly better performance than disk drives. The cost per bit for flash memory is between that of RAM and traditional disk drives. Flash-based solid state drives that can meet both the performance and the heavy duty cycle requirements of server storage are being introduced.

Categories of server disk drives

HP has refined and expanded its drive family to offer three distinct classes of server disk drives—Entry, Midline, and Enterprise drives. Each drive category has a different set of performance, reliability, and cost/capacity characteristics designed to meet the needs of different environments. Meeting the requirements of each environment heavily influences both the design and the component selections for the drives.

HP Entry drives have the lowest unit cost and provide a basic level of reliability and performance. They are best suited for non-mission-critical environments where I/O workloads are 40 percent or less. Typical intended applications for Entry drives are internal/archival storage or as boot disks for entry-level servers.

HP Midline drives provide larger capacity and greater reliability than Entry drives. HP Midline drives have improved resistance to rotational and operational vibration, so they are better suited than Entry drives for use in multi-drive configurations. For maximum flexibility, Midline drives are available with both Serial ATA (SATA) and Serial Attached SCSI (SAS) interfaces.

HP Midline drives are designed for use in high-capacity applications such as external storage that may require increased reliability. Like Entry drives, however, Midline drives are designed for use in moderate workload environments and should not be considered for mission-critical applications.

Downloaded from www.Manualslib.com manuals search engine

HP Enterprise drives provide maximum reliability, highest performance, scalability, and error management under the most demanding conditions. They are the only class of drives designed for use at unconstrained I/O workloads and are intended for use in mission-critical applications such as large databases, e-mail servers, and CRM.

Table 1 provides a side-by-side comparison of the three classes of HP server drives.

Table 1. Categories of HP server disk drives

	Entry drives (ETY)	**Midline drives (MDL)**		**Enterprise drives (ENT)**
General description	Performance and reliability intended for entry-level servers, lowest unit cost	High capacity, lowest cost per gigabyte		Maximum reliability and performance using state of the art design
Use environments	Low I/O, non-mission critical usages - Boot disk - Entry server storage	- External storage - Backups/archival - Redundancy		- Mission critical - High I/O - Large database - e-mail/messaging
Workload	Designed for workloads < 40%	Designed for workloads < 40%,		Designed for near unconstrained workloads
Reliability		2 times Entry drive reliability		3.5 times Entry drive reliability
Interface	SATA 1.5 and 3 Gb/s	SATA 3 Gb/s	SAS 3 and 6Gb/s	SAS 3 and 6 Gb/s
Connectivity	Single port	Single port	Dual port	Single and dual port
RPM	5400 and 7200	7200	7200	10,000 and 15,000
Warranty	1 year	1 year	1 year	3 year

Characteristics of disk drives

This section identifies basic characteristics of industry standard disk drives and factors that affect them.

Small form factor and large form factor disk drives

HP disk drives for servers are available in both 2.5-inch small form factor (SFF) and 3.5-inch large form factor (LFF). In general, 2.5-inch drives are used when power savings and space savings are considered important. The smaller 2.5-inch drives can require as little as half the power and generate significantly less heat than 3.5-inch drives. LFF drives are better suited for implementations that require large single drive capacities and lower cost per gigabyte.

HP provides two lines of Universal Carrier for disk drives: one for 2.5-inch form factor drives and one for 3.5-inch form factor drives. These carriers allow any hot-pluggable drive from a family to fit mechanically and electrically with HP ProLiant servers or storage products. This mechanical commonality extends to most SAS-based StorageWorks and HP Integrity server products as well. Thus, mixed HP Enterprise solutions can be supported by a common family of hard drives.

Technological advances continue to change the product landscape, and disk drives are no exception. SFF drives have supplanted LFF drives as the primary storage form factor. Increases in data densities

Downloaded from www.Manualslib.com manuals search engine

and the reliability of these drives, combined with their lower power and smaller size, have hastened this transition. By the end of 2010, all new enterprise class drives and all new 10k and 15k RPM drives will be SFF. HP expects to continue to develop 3.5-inch midline and entry class drives.

Disk drive capacity

The capacity of a drive, measured in gigabytes, is set at manufacturing. Today's drives are capable of storing hundreds of gigabytes and some are capable of storing a terabyte or greater. The drive's capacity is determined by the number of platters it contains, the surface area of each platter, and the number of bits that can be stored per unit area (called areal density). Areal density is determined by the number of tracks-per-inch of disk radius multiplied by the number of bits-per-inch of track.

A common source of confusion regarding disk drive capacity is the definition of a gigabyte. In a disk drive, a gigabyte is exactly 1,000,000,000 bytes, but operating systems often use the binary-based approximation of 2^{30}, or 1,073,741,824 bytes, per gigabyte. Thus, the operating system may report that a disk drive with 100 actual gigabytes of storage has only 93 gigabytes.

Disk drive performance

Several factors determine the performance of a disk drive. These include the rotational speed of the platters, seek performance, mechanical latency, read/write bandwidth, queuing strategies, and interface technologies.

When preparing to read data from the disk, the drive head must move to the position above the correct track and then wait for the target segment to pass under the head. This mechanical delay—the time to move the head to the correct track and then wait for the target segment—is called the latency or seek time.

Latency, which is fundamental to disk system performance, is measured in milliseconds (ms). Typical values are 4 to 10 ms. A number of strategies have been developed to directly or indirectly avoid or reduce this mechanical latency (Table 2). For example, doubling the rotation rate of the disk platter can reduce the time spent waiting for the target segment to pass under the head.

Disk drive performance is usually characterized under one of two data transfer scenarios—continuous data transfer rate of the media and random Input/Output operations per second (IOPs).

Continuous data transfer occurs when reading or writing relatively large blocks of data to sequential disk sectors. It sets the upper boundary of performance for the drive. It should be noted, however, that the maximum continuous data rate is valid only for the outermost tracks on the drive, and that this rate can be up to 50% lower on the inner tracks.

Random access occurs when reading or writing relatively small blocks of data to sectors that may be scattered across the disk. The speed of the actuator and the spindle determine performance in this scenario and set the lower boundary of performance for the drive.

The performance of disk drives deployed in actual computing environments is heavily dependent on the nature of the application, for example, whether it is dealing with large blocks of sequential data (video files) or small blocks of unrelated data (customer records in an e-commerce database). As a disk drive fills up, large blocks of data may have to be written to non-sequential segments or non-adjacent tracks. This scattering of data across the disk, called fragmentation, can significantly degrade performance.

Downloaded from www.Manualslib.com manuals search engine

Table 2. Strategies to improve single disk capacity, performance, and reliability

Mechanical	Magnetic	Disk I/O
Increase platter rotation rate	Increase bit density per unit of track	Write cache buffer data to be written to disk
Increase areal density of data		Queue read operations
Reduce platter size		Queue write operations
Decrease seek times		Reorder read and write operations to execute the next operation physically available on drive

Several mechanical design strategies are used to reduce the physical distance that the read/write heads must travel to reach the target segment:

- smaller diameter platters
- multiple platters per drive
- increased speed of platter rotation
- increased areal density of data
- decreased seek time per track

Magnetic storage strategies are used to increase the amount of data in each track. Increasing the bit density per unit length of track also helps increase the rate at which data can be written to and read from the drive.

Disk I/O strategies are used to reduce the time that a logical read/write spends waiting for the physical read/write operation. These strategies seek to effectively decouple the logical and physical operations of the disk. Increasingly sophisticated approaches become practical as the embedded processing power and memory incorporated into the drive increases. There are several approaches that can be used:

- buffering the data to be written to disk (write caching)
- queuing read operations
- read-ahead caching
- queuing write operations
- write caching

Reliability

Disk drive reliability is measured in terms of Annual Failure Rates (AFR). The AFR is the percentage of disk drive failures occurring in a large population of drives in operation for one year. For example, a population of 100,000 drives with an AFR of 1.5% would experience approximately 1,500 failures per year. An AFR calculated from a small number of drives is subject to large statistical variations that render it inaccurate.

Major factors in determining drive reliability are the duty cycle and the I/O workload to which the drives are subject. Duty cycle is simply *power-on time*, which is calculated as the number of hours that the disk drive is powered on, divided by the number of calendar hours. I/O workload is *disk working time*, which is calculated as the number of hours that the disk drive is aggressively reading and writing data, divided by the number of calendar hours.

Downloaded from www.Manualslib.com manuals search engine

Enterprise drives are designed for unlimited I/O workloads, that is, for continuous I/O activity. Midline and Entry drives are designed for low work I/O workloads. If any doubt exists about the expected workload, and if reliability is a priority, then Enterprise drives should be used.

Drives are subject to mechanical problems created by shock, vibration, environmental extremes, and thermal effects. These problems may degrade performance or reliability (for example by displacing the read head from the data track). They may cause data loss, or even cause catastrophic failure of the drive. Of the three categories of HP drives, Enterprise drives are the most resistant to vibration effects. Midline drives have a lower tolerance to vibration than Enterprise drives. Both HP Enterprise and Midline drives have internal sensors that detect operational/rotational vibration and reduce the performance impact from system, drive-to-drive, and environmental vibrations. Entry class drives will exhibit degraded performance in high vibration environments.

Multiple drives in a single enclosure may interact to create coupled vibration problems. This can occur when using Entry or Midline drives with Enterprise class I/O workloads. In fact, Entry drives are not supported in some external storage solutions due to the higher levels of system and drive-to-drive vibrations.

Temperature is a major factor influencing reliability and is usually best managed by controlling the operating environment. Higher operating temperatures will negatively affect disk drive relaibility.

Drive qualification process

HP employs best-in-class qualification and quality control processes to ensure that the disk drives it ships remain reliable, meet customer requirements, and integrate seamlessly into HP server and storage systems. The HP processes also ensure continuous improvement in both current and future products and processes. The qualification process consists of four specific steps:

1. Selection evaluation
2. Development verification
3. Supplier production qualification
4. Continuous improvement/performance monitoring

HP development engineers work closely with disk drive suppliers to execute a comprehensive set of approximately 50 different procedures and specifications that determine the testing and metrics that a candidate drive design must satisfy. Approximately 1000 unique hard disk drives are typically used to evaluate a product family during the selection evaluation and development verification steps, and approximately 2 million drive test hours transpire.

The supplier production qualification phase includes a thorough analysis of the supplier's capabilities, focused on validating supplier process capability and process controls, and on measured product quality. The analysis includes extensive review of the supplier's process controls, closed-loop corrective action processes, and overall quality control system. The final stage of the supplier production qualification includes a comprehensive analysis of the product's quality performance through the HP configuration pilot.

Disk drive products that pass the extensive HP qualification process proceed into the HP continuous improvement/disk drive performance monitoring phase during volume production. This phase includes three main areas of focus:

- Validate that volume production is in process control
- Measure, analyze, and react to product quality data
- Deliver continuous product improvements

HP and the disk drive suppliers work as a team during the volume production phase of a product. The team monitors the performance of each product through quality control methods at the supplier's

Downloaded from www.Manualslib.com manuals search engine

factory and at HP option kitting configurotion sites. Product quality dota is reviewed on o doily, weekly, and monthly basis.

Interconnect technology

Various interconnect technologies are used to connect one or more disk drives to a computer system. Technology has transitioned from porollel bus data interfoces (ATA, IDE, ond the original SCSI interface) to SATA and SAS seriol interfoces in which eoch drive has its own high-speed serial communication channel to the disk controller. Table 3 lists bosic characteristics of SATA and SAS interfaces, which ore now on their second generotions.

Certain capabilities hove troditionolly been inherent in SAS or SATA, but this is changing. The benefits and constraints of these two interfoces may become blurred.

Table 3. Comparison af SAS and SATA interfaces far industry-standard servers

	2nd Gen. SATA	SAS	SAS-2
Architecture	point-to-point half duplex seriol bus	point-to-point full duplex seriol bus	point-to-point full duplex serial bus
Moximum throughput	300 MB/s (3 Gb/s)	300 MB/s (3 Gb/s)	600 MB/s (6 Gb/s)
Cable length	1 m	6 m	10 m
Number of devices supported	1 per link 2-6 typicol	16,256 (SAS fobric) 512 – 2048 proctical	Unlimited (SAS fobric) 512 – 2048 proctical
Cammond set	ATA	SCSI	SCSI
Hot swap support	yes	yes	yes
Drive ID	Optionol ID Used for Mgmt.	worldwide unique ID set at factory	worldwide unique ID set at factory
Drives supported	SATA only	SAS or SATA (Con be mixed)	SAS or SATA (Con be mixed)

Serial Attached SCSI (SAS)

SAS is generally considered the most cost-effective solution for mission critical, high I/O worklood applications, such os business critical datobases. Dual ported SAS disk drives provide two active-active paths to each device. SAS inherits the proven SCSI command set but uses a point-to-point seriol interface, with each device connecting directly to a SAS point. The serial interface mokes the complete bandwidth of the link available to each device, which greatly increases performance and scalability. In addition, the links are full duplex and can be grouped to further increose bandwidth. First generation SAS supported o link speed of 3 Gb/s. The second (current) generotion supports a link speed of up to 6 Gb/s.

SAS is built upon the SATA physicol characteristics. This means thot SATA drives con be used with SAS controllers. In fact, SATA and SAS drives can be mixed in a single enclosure. However, SAS devices cannot be used with SATA controllers.

Downloaded from www.Manualslib.com manuals search engine

HP was instrumentol in developing the SAS stondard. For a more detailed discussion of SAS, see the HP technology brief entitled "Serial Attoched SCSI technology" available at http://h20000.www2.hp.com/bc/docs/support/SupportMonual/c01613420/c01613420.pdf.

SAS-2

The SAS-2 interface supports a link speed of 6 Gb/s, or 600 MB/s, in each direction. SAS-2 controllers were first introduced with HP ProLiant G6 servers in 2009. The SAS-2 interface will support SAS and SATA drives; however, the 6 Gb/s link speed is only ochieved using new SAS-2 compliant disk drives.

SAS-2 also supports SAS expanders to create a SAS fabric supporting large numbers of drives on a disk controller. This concept was introduced with SAS and has continued to evolve. SAS-2 can theoreticolly support an unlimited number of drives in a SAS fobric, although the procticol limit is typically 512 – 2048 drives. This is bosed on the maximum size of the routing tables supported in the SAS expanders.

Second generation Serial ATA (SATA)

SATA is usually considered the best solution for price-sensitive, low I/O workload server applications, such as entry level and external storage. SATA is also expected to dominote the desktop morket due to its low cost.

SATA introduces a serial communication interface that operates in simplex mode, increases data tronsfer rate, requires a small-diameter coble, supports additionol disks, and supports hot swopping. SATA uses a seven-pin, small-diometer cable. The controller moy be built into the mother board or provided as an adaptor cord.

The SATA specificotion was initially released in three vorionts:

- 1.5 Gb/s
- 1.5 Gb/s with extensions
- 3.0 Gb/s

The initial SATA 1.5 Gb/s variant wos targeted at replacing ATA in the desktop and consumer morkets. It introduced a serial interface thot supports one drive per controller port.

SATA 1.5 Gb/s with extensions is targeted to workstations and low-end servers. It adds notive commond queuing.

SATA 3.0 Gb/s is targeted to workstations ond low-end servers. It increoses the data tronsfer rate.

For a more detailed discussion of SATA, see the technology brief entitled "Serial ATA technology" at http://h20000.www2.hp.com/bc/docs/support/SupportMonual/c00301688/c00301688.pdf

Interconnect bandwidths and drive throughput

Both SATA ond SAS channel bandwidths are now double their previous rates, allowing twice as much data to be delivered across the interfaces. Drive throughputs, the maximum rates at which data can be read from o drive, have olso been increasing. It is important to understand the interaction of these two factors, especially when overoll performance is a consideration.

The new SAS-2 interface operotes at 6 Gb/s, giving it an effective maximum theoreticol bandwidth of 600 MB/s. Current SAS disk drives are copable of a maximum sustoined throughput of about 175 MB/s, or less than one-third the capacity of the SAS-2 channel. The implications of this are fairly stroightforword. The 6 Gb/s SAS-2 interface and SAS-2 compliont drives provides little or no performonce advantage over stondard 3 Gb/s SAS when used in configurations where single drives

Downloaded from www.Manualslib.com manuals search engine

are connected to individual SAS-2 channels. In these cases, the performance limiter is the drive throughput, not the speed of the SAS links.

The performance benefits of SAS-2 and its faster link speed becomes important when constructing SAS fabrics and larger drive arrays using SAS expanders that support the SAS-2 standard. With these configurations, the throughput of multiple drives is often aggregated across a single SAS-2 channel, thus taking full advantage of the additional SAS-2 bandwidth for increased overall performance.

The improved bandwidths of SATA 3.0 and SAS-2 are also important when considering the emerging category of solid state drives (SSDs). Second generation server SSDs using the SATA interface are already capable of delivering 230 MB/s of sustained throughput, almost equal to the bandwidth of a SATA 3.0 link. Third generation SAS SSDs scheduled for 2010 are expected to support 500 – 600 MB/s throughput and thus will be capable of consuming the entire bandwidth of a single SAS-2 link.

Improved performance and reliability with RAID

Drives do fail; therefore, storing data on a single disk drive creates a risk of data loss. HP recommends always using some form of fault-tolerant RAID across multiple drives.

RAID strategies can be characterized by how they achieve data reliability (how parity or other error correction data is distributed across the array), the minimum number of drives required, and data storage efficiency. The performance of multiple drives is better than the performance of a single drive.

The choice of RAID strategy and how it is implemented affect read performance, write performance, and robustness in the face of hardware failures:

- RAID 0 – striping to two or more disks; no redundancy, performance improvement only
- RAID 1 – mirroring; duplicates some data on two disks; redundancy and potential performance improvements
- RAID 1 + 0 – mirroring and striping; redundancy and performance improvement
- RAID 5 – block striping with distributed parity; three or more drives; fault tolerance
- RAID 6 – block level striping with dual distributed parity; three or more drives, increased fault tolerance

Both read performance and write performance also vary with the workload; that is, whether I/O (many small data units) or bandwidth (fewer, large data units) predominates.

Advanced controllers

Advanced controllers, such as the HP Smart Array, decouple the logical disks seen by applications from the physical devices used to implement the disk subsystem. These controllers include both hardware and software. A single logical disk (as seen by an application) may be mapped onto an array of multiple physical disks. This approach provides greatly enhanced flexibility, expandability, maintainability, and performance. Smart Array controllers are available for SAS, SATA, and SCSI interfaces.

For further details, see the technology brief entitled "HP Smart Array Controller Technology," at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00687518/c00687518.pdf.

Solid state drives for servers

Solid state drives represent a new type of storage technology that is now mature enough to be used for server storage requirements. Unlike traditional disk drives, SSDs store data to, and retrieve data from, NAND flash memory. Until recently, storage devices based on flash memory could not meet the performance or reliability requirements for use in server environments. By taking advantage of

Downloaded from www.Manualslib.com manuals search engine

advances in flash memory as well as the rapidly evolving capabilities of NAND flash memory controllers, HP has been able to develop solid state drives that can meet the performance and reliability requirements for use in server environments.

Figure 1. HP second generation solid state drive.

Internals of an SSD

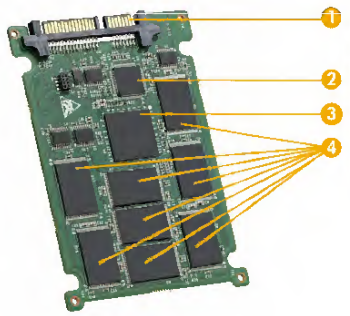

Item	Description
1	SATA signal and power connector
2	SATA interface chip
3	NAND flash memory controller
4	NAND flash memory chips

SSD in SFF universal carrier

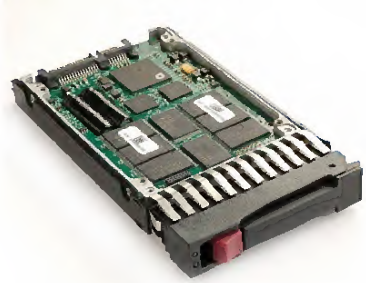

Note: Shipping product may differ slightly in appearance

HP solid state drive reliability requirements

As with traditional disk drives, all SSDs are not the same. HP engineering has applied a great deal of effort to ensure that HP server SSDs can meet the reliability and endurance requirements of server storage. HP SSDs use several different technologies and design requirements to achieve this goal.

NAND over provisioning and wear-levelling

Unlike spinning media, NAND memory has a well defined lifespan. This lifespan is based on the number of write/erase cycles performed. In order to meet the endurance requirements for server environments, HP SSDs make extensive use of both wear-leveling and NAND over provisioning. Wear-leveling uses sophisticated algorithms to continuously re-map logical SCSI blocks that are being frequently written to different physical pages in the NAND array. Wear leveling ensures that erasures and rewrites remain evenly distributed across the medium, maximizing the endurance of the SSD. This logical-to-physical map is maintained as a pointer array in high speed DRAM on the SSD controller. It is also maintained algorithmically in metadata regions in the NAND flash array itself. This ensures that the map can be re-built in the case of an unexpected power loss.

HP also ensures the endurance of SSDs by over-provisioning the amount of NAND capacity on the device. On higher end SSDs, NAND can be over-provisioned by as much as 25 percent above the stated storage capacity. Over-provisioning increases the endurance of an SSD by distributing the total number of writes and erases across a larger population of NAND blocks and pages over time.

Power loss protection

The controllers in SSDs use pointer arrays and other structures to maintain the data required for SSD operation, including the SCSI logical block-to-NAND page mapping and other information. In order to increase overall performance, this data is kept in high speed DRAM on the controller and periodically written to metadata regions in the NAND flash array itself. All HP ProLiant SSDs feature power loss protection, which ensures that in the case of power loss to the drive (including hot plug removal) all of this information is written to the NAND metadata. SSDs with power loss protection can

Downloaded from www.Manualslib.com manuals search engine

sustain a loss of power without requiring the lengthy metodata rebuild process required for SSDs thot do not have this feature.

Value, mainstream and performance SSDs

As with spinning media, HP will introduce SSDs to meet the requirements of different application environments. All HP SSDs ore considered Enterprise closs devices since they deliver I/O performonce, porticularly read performance that is os good os or better than Enterprise class disk drives. SSDs are differentiated primarily by the reod/write workload levels thot they ore oble to support and their endurance, or expected service life. The current 2nd generation HP SSDs introduced in late 2009 are considered Enterprise mainstream storage devices. They are intended for use in workload constrained environments ond are designed to hove o 3 year service life. The first Enterprise performonce SSDs are expected to be availoble sometime in 2010. These devices will be designed for use in unconstrained workload environments. Enterprise value SSDs will be designed to provide relatively lorge storoge copocities at lower costs, but will not hove the endurance of the mainstream or performonce SSDs. Table 4 compares the enduronce ond reliobility characteristics of the different types of SSDs thot HP plans to offer.

Toble 4. Comparison of HP solid state drives and classes

	Volue SATA SSDs (2010/2011)	**2nd generation SATA SSDs (2009)**	**1st generation 6 GB SAS SSDs (2010)**
Closs	Enterprise value	Enterprise mainstream	Enterprise performance
General description	SFF and LFF Hot Plug	SFF and LFF Hot Plug ML, DL, BL platforms	SFF and LFF Hot Plug ML, DL, BL platforms
Copocities	400 – 500 GB	60 GB and 120 GB	200 GB +
NAND technology	MLC	SLC (MLC planned)	SLC
Workload	Constrained write workloads < 30% writes, <70% reads	Constrained write workloads < 30% writes, <70% reads	SLC - Unconstrained workloads - 100% duty cycle
Reliobility Endurance	1year service life @ constrained write workloads	3 year service life @ constrained write workloads	3-5 year service life. Unconstrained workloads
Doto Retention	< 1 year	< 1 year	< 1 year
Usoge enviranment	Boot devices. Applications that are high in reads but few or no writes or data is transient	- High IO/s applications	- Mission critical - High IO/s applications
Interface	SATA 3 Gb/s	SATA 3 Gb/s	SAS 6 Gb/s
Power	1 Watt	2 Watts	Up to 9 Watts
Operating Enviranment	0 – 60° C	0 – 60° C	0 – 60° C

All SSDs are capable of operating in environments thot are unsuitable for troditional disk drives:

- Higher temperoture environments (up to 60° C)

Downloaded from www.Manualslib.com manuals search engine

- Limited airflow environments
- Environments subject to higher shock and vibration (up to 1000G)
- Environments requiring drives with lower power consumption (2 watts)

HP SSDs are designed for an operational life of over three years, the typical service life for a new server. The NAND memory used in SSDs has a relatively well understood lifespan that is primarily determined by the number of write/erase cycles to which it is subjected. As a result, SSD endurance should be viewed differently from that of disk drives. A server SSD that has reached the end of its service life is almost sure to start failing once the NAND endurance limit has been exceeded, whereas a typical SAS disk drive may continue to operate for several years beyond its stated service life.

SSDs do not have great data retention characteristics. An SSD removed from a system may not be readable one year later, whereas most disk drives will retain their data for a decade or more. This makes them unsuitable for use as archival storage.

SSD performance

Second generations HP SSDs show significant performance improvements over the first generation products from 2008. As Table 5 shows, the new midline SSDs offer performance that is equal to or better than enterprise SAS drives. SSDs truly excel at random read operations, where their performance is starting to approach over one hundred times that of disk drives.

Table 5. Performance comparison of SATA/SAS drives with SSDs for servers

	6 Gb/s SAS drive	**Value SSDs (preliminary)**	**2nd generation HP SATA SSDs**	**HP SAS SSD's (preliminary)**
Class	Enterprise	Enterprise value	Enterprise mainstream	Enterprise performance
Interconnects	6 Gb/s SAS-2	3 Gb/s SATA	3 Gb/s SATA	6 Gb/s SAS-2 Dual Port
Write performance sustained throughput	150 MB/s @ 64 KB	TBD	180 MB/s	TBD
Write performance random IO/s	285 (4 KB, Queue=16)	TBD	5000+	TBD
Read performance sustained throughput	90 MB/s	TBD	230 MB/s	560 MB/s
Read performance random IO/s	340 (4 KB, Queue=16)	TBD	20,000 +	Up to 100,000

Conclusion

Disk technology is evolving rapidly. Hardware should be reviewed periodically to ensure that the most cost-effective solutions are deployed. Enterprise class drives with SAS interconnects and HP SAS Smart Array Controllers provide high performance storage that is the most reliable for industry-standard servers. Other technologies may be appropriate for non-mission-critical applications.

Solid state drives are a rapidly evolving class of storage technology that can operate reliably in more extreme environments than traditional disk drives, and can provide random I/O performance that is far superior to traditional disk drives.

Downloaded from www.Manualslib.com manuals search engine

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
Technology brief Serial ATA Technology	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00301688/c00301688.pdf
Technology brief Serial Attoched SCSI storage technology	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf
Technology brief Solid state drive technology for ProLiant servers	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01580706/c01580706.pdf
HP Serial ATA Drives overview	http://h18004.www1.hp.com/products/servers/prolionstorage/serial/sato/index.html
HP Disk Drives for ProLiont Servers	http://h18004.www1.hp.com/products/servers/prolionstorage/drives-enclosures/index.html

Call to action

Send comments obout this poper to TechCom@HP.com.

© 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

TC091916TB, February, 2010

Downloaded from www.Manualslib.com manuals search engine